arXiv:1607.06051v1 [stat.ME] 20 Jul 2016

A Bayesian Model for Aggregating Rank Data with Covariates

Dingdong Yi, Xinran Li and Jun S. Liu*

July 21, 2016

Abstract

Rank aggregation is the combining of ranking results from different sources to generate a "better" ranking list. Existing rank aggregation methods utilize only rank data. In some real applications, however, rank data come with covariates information as well. Moreover, some rankers may be more informative and reliable than others. Therefore, it is desirable for a rank aggregation method to incorporate covariates information and to distinguish high quality rankers from spam rankers. Here, we propose the Bayesian aggregation of rank-data with covariates (BARC) and its weighted version, BARCW, to achieve the aforementioned goals. Both BARC and BARCW employ latent variable models—which can incorporate covariates information—for ranked entities. Furthermore, BARCW quantitatively measures the difference in qualities of rankers. We generate an aggregated ranking and its probability interval based on posterior distributions of latent variables derived from a parameter-expanded Gibbs sampler. Simulation studies show that our methods outperform competing methods in a variety of scenarios. Real-data applications also validate our estimation of reliability parameters and effects of covaraiates.

Keywords: Latent variable model; Rank aggregation; PX-DA; Spam ranker detection

*Dingdong Yi (E-mail: ddyi@fas.harvard.edu) and Xinran Li (E-mail: xinranli@fas.harvard.edu) are Doctoral Candidates, Department of Statistics, Harvard University, Cambridge, MA 02138. Jun S. Liu is Professor of Statistics, Department of Statistics, Harvard University, Cambridge, MA 02138 (E-mail: jliu@stat.harvard.edu).

1 Introduction

Combining ranking results from different sources is a common problem. People expect the aggregated rank to provide "better" information than each individual rank. The beginnings of rank aggregation can be traced back to studies of social choice theory and political elections in the 18th century (de Borda, 1781). Thurstone (1927) proposed a model based approach to analyze pairwise comparison data in psychology by introducing latent variables. With the rise of internet search engines, it has become necessary and challenging to compute an aggregated "consensus" ranking list to eliminate the bias of individual ranking lists, which usually contain thousands of entities. Dwork et al. (2001) proposed simple Markov Chain based methods in this context. Lin and Ding (2009) applied a stochastic optimization-based method to aggregate a top-d list from multiple platforms of biological study. Deng et al. (2014) proposed a Bayesian top-d list aggregation method to address differences in the reliability of the base rankers.

Existing rank aggregation methods utilize only score or rank information, whereas in some real applications, rank data also come with covariates. Covariates provide extra information about ranking lists and rankers, and can help aggregate some highly incomplete ranking lists, which existing methods cannot aggregate. Below are two motivating real data examples:

Example 1 (NFL Quarterback Ranking) *During the NFL season, experts from different websites, such as* `ESPN.com` *and* `nfl.com`*, provide weekly ranking lists of players by position. Table 1 shows the ranking lists of NFL starting quarterbacks, ranked in week 12 of the 2014 NFL season by 13 experts. The ranking lists can help football fans better predict the performance of the quarterbacks in the coming week and even place bets in online fantasy sports games. After collecting ranking lists from rankers, these websites mostly aggregate them using arithmetic means. We notice that NFL players' summary statistics are also available online. Table 2 is an example showing these ranked quarterbacks' statistics prior to week 12 of the 2014 NFL season. In their analysis, the rankers use these summary statistics in addition to watching football games, but each may use the summary statistics in a different way.*

Table 1: Ranking lists of NFL starting quarterbacks from 13 different experts, as of week 12 in the 2014 season.

No.	Player	τ_1	τ_2	τ_3	τ_4	τ_5	τ_6	τ_7	τ_8	τ_9	τ_{10}	τ_{11}	τ_{12}	τ_{13}
1	Andrew Luck	1	1	1	3	3	1	1	1	1	1	1	1	1
2	Aaron Rodgers	2	3	4	2	1	2	3	3	2	2	3	4	3
3	Peyton Manning	3	2	5	4	2	3	2	2	3	4	4	2	2
4	Tom Brady	4	7	3	5	4	5	4	6	4	3	6	8	4
5	Tony Romo	9	5	6	1	5	4	5	4	5	5	7	6	6
6	Drew Brees	10	4	2	8	9	7	7	5	7	6	2	3	5
7	Ben Roethlisberger	6	8	7	7	7	6	6	10	6	7	5	7	7
8	Ryan Tannehill	5	6	13	6	11	8	8	7	9	9	8	5	8
9	Matthew Stafford	8	9	11	13	8	9	9	8	8	8	9	9	9
10	Mark Sanchez	22	10	9	9	16	10	10	9	10	10	12	12	12
11	Russell Wilson	12	13	17	10	10	12	11	12	11	12	11	14	15
12	Philip Rivers	7	14	15	20	6	17	17	11	16	15	14	10	10
13	Cam Newton	18	12	8	17	19	11	14	14	14	16	21	13	14
14	Eli Manning	17	–	18	19	14	19	12	13	12	13	16	23	11
15	Matt Ryan	21	17	19	15	20	15	15	15	13	11	20	21	13
16	Andy Dalton	15	–	14	–	17	14	16	20	15	14	19	22	16
17	Alex Smith	16	11	21	16	18	18	18	16	20	21	13	11	17
18	Colin Kaepernick	11	16	16	11	12	16	21	17	19	18	22	16	21
19	Joe Flacco	24	15	12	14	24	13	13	18	18	20	15	15	19
20	Jay Culter	13	18	10	12	13	21	19	19	17	17	23	20	18
21	Josh McCown	14	19	22	18	15	22	22	21	21	19	18	17	23
22	Drew Stanton	20	20	–	22	22	20	20	23	22	22	10	19	20
23	Teddy Bridgewater	23	21	20	21	23	23	23	22	23	24	17	18	22
24	Brian Hoyer	19	–	–	–	21	24	24	24	24	23	24	24	24

Source: http://fantasy.nfl.com/research/rankings, http://www.fantasypros.com/nfl/rankings/qb.php

Example 2 (Orthodontics treatment evaluation ranking) *In 2009, 69 orthodontics experts were invited by Peking University's School of Stomatology to evaluate the pre-treatment and post-treatment conditions of 108 medical cases. In order to make the evaluation easier for experts, cases were divided into 9 equal groups. For each group of the cases, each expert evaluated the conditions of cases and provided a within-group ranking as a result. In the meantime, using each case's plaster model, cephalometric radiograph, and photograph, the School of Stomatology located key points and measured their distances and angles, which are considered to be relevant features for diagnosis. Thus, similar to the quarterback ranking example, the data are composed of objective covariates and subjective rankings. Table 3 shows 20 of the 69 ranking lists for two groups. Because there is no comparison between groups, existing methods cannot provide us with a full aggregated ranking list.*

We propose the Bayesian aggregation of rank-data with covariates (BARC) and its weighted version (BARCW) to tackle rank aggregation problem where relevant covariates are also observed. Similar to Thurstone (1927), we introduce an underlying latent value

Table 2: Relevant statistics of the ranked quarterbacks, prior to week 12 of the 2014 NFL season.

No.	Player	G	Pct	Att/G	Avg	Yds/G	TD%	Int%	RAtt/G	RAvg	RYds/G	R1st%
1	Andrew Luck	11	63.40	42.20	7.80	331.00	6.30	2.20	4.20	4.20	17.50	30.40
2	Aaron Rodgers	11	66.70	31.10	8.60	268.80	8.80	0.90	2.50	6.40	16.20	50.00
3	Peyton Manning	11	68.10	40.20	8.00	323.50	7.70	2.00	1.50	-0.50	-0.70	0.00
4	Tom Brady	11	65.00	37.90	7.20	272.50	6.20	1.40	1.70	0.70	1.30	21.10
5	Tony Romo	10	68.80	29.50	8.50	251.90	7.50	2.00	1.50	2.50	3.70	20.00
6	Drew Brees	11	70.30	42.00	7.60	317.40	4.80	2.40	1.70	2.80	4.90	26.30
7	Ben Roethlisberger	11	68.30	37.50	7.90	297.30	5.80	1.50	1.90	1.10	2.10	19.00
8	Ryan Tannehill	11	66.10	35.40	6.60	234.70	5.10	2.10	3.70	6.70	25.10	36.60
9	Matthew Stafford	11	58.80	37.70	7.10	267.50	3.10	2.40	2.80	2.00	5.60	16.10
10	Mark Sanchez	4	62.30	36.50	8.10	296.80	4.80	4.10	3.50	0.60	2.00	7.10
11	Russell Wilson	11	63.60	28.50	7.10	202.70	4.50	1.60	7.60	7.70	58.50	45.20
12	Philip Rivers	11	68.30	33.00	7.80	257.70	6.10	2.50	2.50	2.50	6.40	25.00
13	Cam Newton	10	58.60	33.30	7.20	239.20	3.60	3.00	6.40	4.60	29.30	37.50
14	Eli Manning	11	62.30	36.90	7.00	257.50	5.20	3.00	0.80	3.80	3.10	33.30
15	Matt Ryan	11	65.10	38.50	7.20	278.70	4.50	2.10	1.60	4.30	7.10	33.30
16	Andy Dalton	11	62.40	30.70	7.10	219.40	3.60	3.00	3.80	2.50	9.50	33.30
17	Alex Smith	11	65.10	29.70	6.80	201.00	4.00	1.20	3.20	5.50	17.40	25.70
18	Colin Kaepernick	11	61.70	31.50	7.50	237.70	4.30	1.70	6.80	4.50	30.50	22.70
19	Joe Flacco	11	63.20	34.10	7.40	251.30	4.80	2.10	2.00	1.70	3.40	45.50
20	Jay Cutler	11	66.80	36.40	7.10	256.80	5.50	3.00	2.90	3.90	11.30	28.10
21	Josh McCown	6	60.40	30.30	7.40	225.00	3.80	4.40	2.70	5.80	15.30	50.00
22	Drew Stanton	6	53.60	25.20	7.10	178.20	3.30	2.00	3.00	2.00	6.00	22.20
23	Teddy Bridgewater	8	60.30	32.80	6.40	211.10	2.30	2.70	3.50	4.60	16.10	32.10
24	Brian Hoyer	11	55.90	33.20	7.80	260.40	3.00	2.20	1.80	0.90	1.50	20.00

Source: http://www.nfl.com/stats

vector for ranked entities to transform a rank aggregation problem into a Bayesian inference of the latent entity vector. We assume that each ranking list is determined by its ranker's value vector, a noisy version of the latent entity vector. We believe that reliable rankers can generate more consistent ranking lists than spam rankers, if there is replication. By assuming differences in the reliability of rankers, which is inversely proportional to the noise level, BARCW learns the rankers' reliability from data. We use a parameter-expanded Gibbs sampler to overcome some difficulties in Bayesian computation. In contrast to existing rank aggregation methods, our method is adaptive to unequal and unknown reliability of input lists, and can incorporate covariates to aggregate incomplete ranking lists robustly.

The rest of this article is organized as follows: Section 2 introduces basic definitions and distance measures of ranking lists; Section 3 reviews some existing rank aggregation methods; Section 4 elaborates on our Bayesian models for rank data with covariates; Section 5 provides details of our Markov Chain Monte Carlo algorithm; Section 6 describes simulation studies under different situations; Section 7 applies our methods, BARC and BARCW, to two real

Table 3: Ranking lists for Group A & B, 2 of the 9 groups in Example 2

	R1	R2	R3	R4	R5	R6	R7	R8	R9	R10	R11	R12	R13	R14	R15	R16	R17	R18	R19	R20
A1	1	3	5	2	4	1	1	2	5	5	10	8	2	4	2	2	3	1	1	5
A2	11	5	10	9	9	12	9	7	11	12	4	7	5	6	5	9	4	10	10	1
A3	6	10	8	11	11	8	11	8	12	9	6	11	12	11	11	12	12	9	11	11
A4	3	2	4	3	1	4	2	10	1	6	8	2	1	1	1	1	2	2	4	2
A5	9	4	7	5	6	6	6	5	3	3	2	5	11	7	9	4	9	8	8	7
A6	10	9	3	6	5	11	5	9	6	7	3	1	6	8	7	5	6	5	5	8
A7	8	8	11	7	12	9	12	11	8	10	7	9	8	12	12	10	7	7	12	9
A8	4	1	1	4	3	2	4	4	2	1	1	6	3	2	6	7	5	4	7	3
A9	2	12	9	8	8	5	7	3	9	8	11	12	7	5	8	6	11	12	2	12
A10	7	11	6	10	10	7	8	6	7	11	9	3	10	9	4	3	1	11	3	6
A11	5	7	2	1	2	3	10	1	10	2	5	4	9	3	3	8	8	6	6	10
A12	12	6	12	12	7	10	3	12	4	4	12	10	4	10	10	11	10	3	9	4
B1	3	2	9	7	5	6	4	8	3	4	3	1	5	2	2	1	2	4	11	8
B2	7	5	5	10	10	10	12	2	8	11	4	6	7	10	11	6	11	11	12	12
B3	1	1	1	5	1	4	3	1	4	1	2	4	4	6	1	2	8	2	1	3
B4	9	7	4	4	4	5	8	9	12	6	10	11	8	3	6	3	3	9	6	10
B5	4	9	8	8	3	9	9	10	6	7	8	5	9	8	5	9	10	8	5	11
B6	2	4	6	1	8	2	2	4	5	10	6	10	6	7	4	5	9	6	3	4
B7	8	8	2	6	2	7	7	3	1	3	9	7	10	4	9	7	6	7	2	2
B8	6	10	3	3	6	1	5	5	10	2	5	3	1	1	7	10	5	1	4	1
B9	12	11	11	12	9	8	6	6	9	9	11	8	11	11	12	12	12	10	10	5
B10	10	6	10	2	7	12	1	11	2	8	7	9	3	12	8	8	7	5	7	6
B11	11	12	12	11	12	11	11	12	11	12	12	12	12	9	10	11	4	12	9	9
B12	5	3	7	9	11	3	10	7	7	5	1	2	2	5	3	4	1	3	8	7

data examples; and Section 8 concludes the article with a short discussion.

2 Notation and definitions

2.1 Ranking list

Let U be a set consisting of all entities, and $|U|$ be the total number of entities in U. We introduce a binary relation '$\succ$' for entities in U, where $i_1 \succ i_2$ if entity i_1 is ranked higher than entity i_2. The binary relation '$\succ$' has the properties: (1) irreflexivity, i.e. not $i \succ i$; (2) transitivity, i.e. if $i_1 \succ i_2$ and $i_2 \succ i_3$, then $i_1 \succ i_3$.

A *ranking list* $\tau = \{\tau[i_1, i_2] = 1_{\{i_1 \succ i_2\}} : i_1, i_2 \in U\}$ is a set of pairwise relations in U. In particular, τ is a *full ranking list* (or simply, a *full list)* if τ identifies all pairwise comparisons in U. When τ is a full ranking list, we can rewrite τ as $\tau = [y_1 \succ y_2 \succ \ldots \succ y_{|U|}]$ for notational simplicity, and further define $\tau(i)$ as the position of entity $i \in U$ (a highly ranked element has a low-numbered position in the list, i.e. $\tau(i_1) < \tau(i_2)$ if $i_1 \succ i_2$). A ranking list τ^P is said to be a *partial list* if any of the $\binom{|U|}{2}$ pairwise comparisons is missing.

If τ contains all elements in a subset $T \subseteq U$, then the projection of τ with respect to T, denoted as $\tau_{|T}$, is a ranking list with respect to T. For example, in our orthodontics data, S_g is the set of cases in each group, and $U = S_1 \cup S_2 \cup \cdots \cup S_9$ is the set of all cases, where $|S_g| = 12$ and $|U| = 108$. For each expert j ($1 \leq j \leq 69$), we observe a partial list τ_j^P, where $\tau_{j|S_g}^P$ is a full list of S_g.

2.2 Distance measures

We focus on the following two popular distance measures between ranking lists.

(1) *Spearman's footrule distance* (Diaconis and Graham, 1977). $F(\sigma, \tau)$ is the sum, over all entities $i \in U$, of absolute difference between ranking positions $\sigma(i)$ and $\tau(i)$ in two lists; that is, $F(\sigma, \tau) = \sum_{i=1}^{|U|} |\sigma(i) - \tau(i)|$. Dividing $F(\sigma, \tau)$ by its maximum value $\frac{|U|^2}{2}$, we have $d_F(\sigma, \tau)$ as a distance between taking a value 0 and 1. The normalized Footrule distance of σ to several full lists $\tau_1, \ldots, \tau_m$ is defined as

$$d_F(\sigma, \tau_1, \ldots, \tau_m) = \frac{1}{m} \sum_{j=1}^{m} d_F(\sigma, \tau_j) \tag{1}$$

(2) *Kendall tau distance* (Kendall, 1938). $K(\sigma, \tau)$ counts the number of all discordant pairs in two lists; that is $K(\sigma, \tau) = \Big|\{(i_1, i_2) | \sigma(i_1) < \sigma(i_2), \text{ but } \tau(i_1) > \tau(i_2)\}\Big|$. Dividing $K(\sigma, \tau)$ by its maximum value $\frac{|U|(|U|-1)}{2}$, we have $d_K(\sigma, \tau)$ as a distance taking a value between 0 and 1. The normalized Kendall tau distance of σ to several full lists $\tau_1, \ldots, \tau_m$ is defined as

$$d_K(\sigma, \tau_1, \ldots, \tau_m) = \frac{1}{m} \sum_{j=1}^{m} d_K(\sigma, \tau_j) \tag{2}$$

3 Overview of Rank Aggregation Methods

3.1 Methods based on summary statistics

Rank aggregation methods based on summary statistics (e.g. average ranking position) are easily understood and widely used. Suppose we have m full ranking lists. Let $\{\tau_j(i)\}_{1 \leq j \leq m}$ be the ranking positions of entity i received from all m rankers. The Borda Count method, proposed by de Borda (1781), aggregates ranks based on their arithmetic mean, $\sum_{j=1}^{m} \tau_j(i)/m$.

Similar aggregation functions are the geometric mean and the median of $\{\tau_j(i)\}_{1\leq j\leq m}$. We refer these three methods as AriM, GeoM and MedR, respectively. These methods perform well when ranking lists are full and of similar quality. But it is difficult for these methods to deal with partial rankings lists and/or when rankers are of different reliability.

3.2 Thurstone's ranking model

Thurstone's model for paired comparisons (Thurstone, 1927) assumes an underlying value vector for the items, $\boldsymbol{\mu} = (\mu_1, \mu_2, \ldots, \mu_{|U|})$, and a latent vector $\boldsymbol{Z} = (Z_1, Z_2, \ldots, Z_{|U|})$ for each ranker. Specifically, it means that the latent vector $\boldsymbol{Z}$ for each ranker follows a multivariate normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Under this model, any pair of underlying values Z_{i_1} and Z_{i_2} follows the bivariate normal distribution

$$(Z_{i_1}, Z_{i_2}) \sim BVN\left(\mu_{i_1}, \mu_{i_2}, \sigma_{i_1}^2, \sigma_{i_2}^2, \rho_{i_1 i_2}\right),$$

where σ_i's are specified as discriminal dispersions and $\rho_{i_1 i_2}$ is the correlation between the discriminal deviations of entity i_1 and i_2. Therefore, the probability that entity i_1 ranks higher than entity i_2 is

$$P(Z_{i_1} > Z_{i_2}) = \Phi\left(\frac{\mu_{i_1} - \mu_{i_2}}{\sqrt{\sigma_{i_1}^2 + \sigma_{i_2}^2 - 2\rho_{i_1 i_2}\sigma_{i_1}\sigma_{i_2}}}\right)$$

where Φ is the cumulative distribution function of the standard normal distribution. Intuitively, when we have m full ranking lists, the observed frequency of pairwise comparison $\frac{1}{m}\sum_{j=1}^{m} 1_{\{\tau_j(i_1)<\tau_j(i_2)\}}$ should approximate $P(Z_{i_1} > Z_{i_2})$ and help measure relative scale difference $\mu_{i_1} - \mu_{i_2}$.

Thurstone's model was later extended to an inference of the whole vector $\boldsymbol{\mu}$ and it has been extensively used in marketing research (e.g., consumer rankings of products) and psychology studies (Lin, 2010). Under the model setting, the number of unknown parameters is as large as $\frac{|U|(|U|+3)}{2}$, so some of them need to be fixed in order to make the parameters identifiable (e.g., let $\mu_1 = 1$ and $\sigma_1^2 = 1$). It is not numerically feasible to find the MLE for these parameters, although various estimation methods have been developed. Maydeu-Olivares (1999) proposed a non-linear least square estimator for Thurstone's model; DeConde et al.

(2006) applied this proposed estimator to aggregate the results of several microarray studies by setting $\rho_i = 0$ and $\sigma_i^2 = 1$ for all i. Yao and Böckenholt (1999) proposed a Bayesian estimation of Thurstonian ranking models based on the Gibbs sampler, in a situation where the length of ranking lists is small (e.g., $4 \leq |U| \leq 10$).

3.3 Stochastic optimization-based and weighted rank aggregation

Depending on the distance measure for ranking lists, optimization-based methods have been proposed to minimize the average distance between a candidate list and each of the input lists, i.e.,

$$\rho = \arg \min_{\sigma \in \mathcal{S}(U)} d(\sigma, \tau_1, \ldots, \tau_m) \tag{3}$$

where $\mathcal{S}(U)$ represents all allowable rankings, and $d(\cdot)$ is either the average Spearman's footrule distance in (1) or the average Kendall tau distance in (2).

When minimizing average Kendall tau distance, ρ is referred to as the *Kemeny optimal aggregation* (Dwork et al., 2001). The ranking obtained by optimizing the Spearman's footrule distance is called the *Footrule optimal aggregation*. The Kemeny optimal aggregation gives the maximum likelihood estimator of the "correct opinion" under a voting model, in which, for each pairwise comparison, each voter would agree with the "correct opinion" with some fixed probability greater than 1/2 (Young, 1988).

Lin and Ding (2009) extended this objective function to the following weighted version,

$$\rho = \arg \min_{\sigma \in \mathcal{S}(U)} d(\sigma, \tau_1, \ldots, \tau_m; w) = \frac{1}{m} \sum_{j=1}^{m} w_i d(\sigma, \tau_j), \tag{4}$$

where $w = (w_1, ..., w_m)$ (subject to $\sum_{j=1}^{m} w_j = 1$ and $w_j \geq 0$) is a weight vector that quantifies prior information on the quality or importance of the individual ranking lists.

In terms of computational complexity, the Kemeny optimal aggregation is a quadratic assignment problem, which is NP-hard; the Footrule optimal aggregation, on the other hand, requires an expensive polynomial algorithm. Therefore, Lin and Ding (2009) used a stochastic search method to optimize (4) by adopting the cross entropy Monte Carlo (CEMC) approach (Rubinstein and Kroese, 2004). The corresponding optimization methods based on these two distance measures are denoted as CEMC_K and CEMC_F.

3.4 Markov Chain Based Methods

Dwork et al. (2001) proposed three Markov Chain based methods (MC_1, MC_2, MC_3) to solve the rank aggregation problem. The basic idea behind these methods is to construct a Markov chain with transition matrix $P = \{p_{i_1 i_2}\}_{i_1, i_2 \in U}$, where $p_{i_1 i_2}$ is the transition probability from entity i_1 to entity i_2, based on the pairwise comparison information from $\{\tau_1, \ldots, \tau_m\}$. For example, the transition rule of MC_2 is:

> If the current state is i_1 then the next state is chosen by first picking a list τ uniformly from all the partial lists $\{\tau_1, \ldots, \tau_m\}$ containing entity i_1 then picking an entity i_2 uniformly from the set $\{i_2 \mid \tau(i_2) \leq \tau(i_1)\}$.

Then, the authors use the stationary distribution of this Markov chain to generate the aggregated ranking list ρ. Explicitly,

$$\rho = \text{sort}(i \in U \text{ by } \pi_i \downarrow),$$

where $\pi = (\pi_1, \ldots, \pi_{|U|})$ satisfies $\pi P = \pi$, and the symbol "$\downarrow$" means that the entities are sorted in descending order.

Markov Chain based methods are simple to implement and fast even when the ranking lists are fairly long (e.g., search engine results). From a statistical point of view, however, the underlying assumptions and interpretations of these methods are not clear.

4 Bayesian models for rank data with covariates

4.1 Full ranking lists without covariates

Suppose we have m full ranking lists $\tau_1, \tau_2, \ldots, \tau_m$ in $U = \{1, 2, \ldots, n\}$. Thurstone (1927) postulated that the ranking outcome τ_j is determined by a latent random variable Z_{ij} ($i = 1, \ldots, n$). The variable Z_{ij} is ranker j's evaluation for the ith entity, such that $Z_{i_1 j} > Z_{i_2 j}$ when $\tau_j(i_1) < \tau_j(i_2)$. In other words, an entity is ranked first in a ranking list if its latent value is the largest for this ranker, etc. With a slight abuse of our notations, for each ranker j, we define $\boldsymbol{Z}_j = (Z_{1j}, \ldots, Z_{nj})'$ as its evaluation for all entities, and $\tau(\boldsymbol{Z}_j) = \text{sort}(i \in U \text{ by } Z_{ij} \downarrow)$

as the ranking list of latent variables. Similar to Thurstone's assumption, we assume that $\boldsymbol{Z}_j$ follows a multivariate normal distribution,

$$\begin{aligned} Z_{ij} &= \mu_i + \epsilon_{ij}, \quad (1 \le i \le n) \\ \tau_j &= \tau(\boldsymbol{Z}_j), \quad (1 \le j \le m) \end{aligned} \tag{5}$$

where $\epsilon_{ij} \sim N(0, \sigma^2)$ independently. Multiplying $(\boldsymbol{\mu}, \sigma)$ by a constant or adding a same constant to μ_is does not influence the observed likelihood, since we only observe ranking lists τ_js. To make the model identifiable, we fix $\sigma^2 = 1$ and restrict $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)'$ to the parameter space $\boldsymbol{\Theta} = \{\boldsymbol{\mu} \in \mathbb{R}^n : \mathbf{1}'\boldsymbol{\mu} = 0\}$.

The model above implies that the τ_j's are independent and identically distributed (i.i.d.) conditioning on $\boldsymbol{\mu}$, so the likelihood function is

$$p(\tau_1, \cdots, \tau_m \mid \boldsymbol{\mu}) = \prod_{j=1}^{m} p(\tau_j \mid \boldsymbol{\mu}) = \prod_{j=1}^{m} \int p(\tau_j \mid \boldsymbol{Z}_j, \boldsymbol{\mu}) p(\boldsymbol{Z}_j \mid \boldsymbol{\mu}) d\boldsymbol{Z}_j,$$

where $p(\tau_j \mid \boldsymbol{Z}_j, \boldsymbol{\mu}) = 1_{\{\tau(\boldsymbol{Z}_j) = \tau_j\}}$. Explicitly, for any possible full ranking list τ' on $U = \{1, 2, \ldots, n\}$, the probability mass function is

$$P(\tau_j = \tau' \mid \boldsymbol{\mu}) = \int_{\tau(\boldsymbol{Z}_j) = \tau'} (2\pi)^{-n/2} e^{-\frac{1}{2}\|\boldsymbol{Z}_j - \boldsymbol{\mu}\|^2} d\boldsymbol{Z}_j.$$

Our goal is to generate the aggregated rank based on an estimate of $\boldsymbol{\mu}$ in model (5). One approach is to use the maximum likelihood estimator (MLE) $\hat{\boldsymbol{\mu}}_m$ that maximizes the log-likelihood function $L_m(\boldsymbol{\mu})$, i.e.,

$$\hat{\boldsymbol{\mu}}_m = arg\max_{\boldsymbol{\mu}} \frac{1}{m} \sum_{j=1}^{m} \log p(\tau_j | \boldsymbol{\mu}).$$

We have the following consistency result, whose proof is deferred to the appendix.

Theorem 1 *Let true parameter value of model (5) be $\boldsymbol{\mu}_0 \in \boldsymbol{\Theta}$, and we observe $\tau_1, \ldots, \tau_m$ generated from the model. Let $\hat{\boldsymbol{\mu}}_m$ be the MLE of $\boldsymbol{\mu}_0$. Then, $\forall \epsilon > 0$ and compact set $K \subset \boldsymbol{\Theta}$, we have $P\left(\{\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}_0\| \ge \epsilon\} \cap \{\hat{\boldsymbol{\mu}}_m \in K\}\right) \to 0$, as n is fixed and $m \to \infty$.*

Alternatively, we can employ a Bayesian procedure to make inference from model (5), which is more convenient to incorporate prior information, to quantify estimation uncertainties, and to implement data augmentation (Tanner and Wong, 1987) and parameter

expansion strategies for efficient MCMC computation. With a reasonable prior, the posterior mean of $\boldsymbol{\mu}$ is also a consistent estimator because of the consistency of the MLE.

With prior $p(\boldsymbol{\mu})$, the posterior distribution of $\boldsymbol{\mu}$ and $\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m$ is

$$p(\boldsymbol{\mu}, \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m \mid \tau_1, \cdots, \tau_m) \propto p(\tau_1, \cdots, \tau_m \mid \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m, \boldsymbol{\mu}) p(\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m \mid \boldsymbol{\mu}) p(\boldsymbol{\mu})$$

The posterior mean of $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)'$ is

$$\bar{\mu}_i = \int \mu_i p(\mu_i \mid \tau_1, \cdots, \tau_m) d\mu_i, \quad (1 \le i \le n)$$

where $p(\mu_i \mid \tau_1, \cdots, \tau_m)$ is the marginal posterior distribution. Based on $\bar{\mu}_i$s, we can generate the aggregated ranking list

$$\rho = \text{sort}(i \in U \text{ by } \bar{\mu}_i \downarrow).$$

Let $\boldsymbol{P}_n = \boldsymbol{I}_n - n^{-1}\mathbf{1}\mathbf{1}'$ denote the projection matrix on $\mathbb{R}^n$ such that the range of projection $\boldsymbol{P}_n$ is $\boldsymbol{\Theta}$. We let the prior for $\boldsymbol{\mu}$, which is restricted on parameter space $\boldsymbol{\Theta}$, be $\mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{P}_n)$. The intuition for choosing this prior is that when $\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{I}_n)$, we have $\boldsymbol{\mu} - \bar{\boldsymbol{\mu}} = \boldsymbol{P}_n \boldsymbol{\mu} \in \boldsymbol{\Theta}$ and $\boldsymbol{\mu} - \bar{\boldsymbol{\mu}} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{P}_n)$, where $\bar{\boldsymbol{\mu}} \overset{\text{def}}{=} \frac{1}{n} \sum_{i=1}^n \mu_i \mathbf{1} = \frac{1}{n} \mathbf{1}\mathbf{1}' \boldsymbol{\mu}$. In computation, it is equivalent to using the prior $\boldsymbol{\mu} \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{I}_n)$ and considering the posterior mean of $\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}$. In other words,

$$p_{\pi_1}(\boldsymbol{\mu} \mid \tau_1, \cdots, \tau_m) = p_{\pi_2}(\boldsymbol{\mu} - \bar{\boldsymbol{\mu}} \mid \tau_1, \cdots, \tau_m)$$

where $\pi_1 \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{P}_n)$ and $\pi_2 \sim \mathcal{N}(\mathbf{0}, \sigma_0^2 \boldsymbol{I}_n)$.

Our estimation is different from typical methods in Thurstone's model which fix one entry of $\boldsymbol{\mu}$ (e.g. $\mu_1 = 0$), as used in Yao and Böckenholt (1999). Under our model setting and prior choice, it is convenient to extend this model using covariate information. We will discuss the sensitivity to hyper-parameter σ_0 in simulation studies.

4.2 Ranking lists with covariates

In real data, the ranking lists can be incomplete, and can also be associated with relevant covariates. By adding a linear combination of covariates to model (5), we have

$$\begin{aligned} \mu_i &= \alpha_i + x_i^T \boldsymbol{\beta}, \\ Z_{ij} &= \mu_i + \epsilon_{ij}, \ \epsilon_{ij} \sim \mathcal{N}(0, 1) \\ \tau_j^* &= \tau(\boldsymbol{Z}_j) \end{aligned} \tag{6}$$

where $\tau(z)$ is the full ranking list of z and $\boldsymbol{\mu} \in \boldsymbol{\Theta} \stackrel{def}{=} \{\boldsymbol{\mu} \in \mathbb{R}^n : \mathbf{1}'\boldsymbol{\mu} = 0\}$. We observe partial ranking lists $\{\tau_j^P\}_{1\le j\le m}$ and covariates $\{x_i\}_{1\le i\le n}$. We denote $\tau^* \sim \tau^P$ if all pairwise comparisons in the full list τ^* agree with the partial list τ^P.

For example, for the quarterback data, the rationale behind this model can be explained as follows. The unobserved variable Z_{ij} represents ranker j's performance evaluation for quarterback i. The expression $\alpha_i + x_i^T\boldsymbol{\beta}$ represents the true performance of the quarterback, but each ranker evaluates it with a personal variation modeled by ϵ_{ij}. The term $x_i^T\boldsymbol{\beta}$ explains the part of their performance that can be explained perfectly by the covariates. The term α_i explains the remaining part of the quarterbacks' performances that is not captured by their summary statistics. In other words, summary statistics do not capture all characteristics in live games.

Model (6) is indeed an over-parameterized version of model (5): τ is invariant if we add a constant vector $\boldsymbol{c}$ to $\boldsymbol{\beta}$ and change α_i to $\alpha_i - x_i^T\boldsymbol{c}$. Although the parameters $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are non-identifiable without additional constraints, we only use the estimation of $\boldsymbol{\mu}$ for rank aggregation. As we did in Section 3.1, we restrict $\boldsymbol{\mu}$ to parameter space $\boldsymbol{\Theta}$, which does not change the ranking of μ_is, to make $\boldsymbol{\mu}$ identifiable.

Note that any prior of $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ implies a prior of the unconstrained $\boldsymbol{\mu}$. To specify the prior $p(\boldsymbol{\alpha}, \boldsymbol{\beta}) = p(\boldsymbol{\alpha})p(\boldsymbol{\beta})$ is computationally easier than to directly specify a prior distribution for the constrained $\boldsymbol{\mu} \in \boldsymbol{\Theta}$. With this prior, the posterior mean of $\boldsymbol{\mu} - \bar{\boldsymbol{\mu}}$ (here $\bar{\boldsymbol{\mu}} \stackrel{\text{def}}{=} \frac{1}{n}\sum_{i=1}^n \mu_i \mathbf{1}$) is our estimator for $\boldsymbol{\mu} \in \boldsymbol{\Theta}$. Essentially, model (6) incorporates the covariates information into the prior. It performs well in finite sample cases where the ranking information may be incomplete but the covariates information is relevant for ranking. In practice, we assume a Gaussian prior $\mathcal{N}(0, \sigma_0^2 I)$ constrained in $[-\lambda, \lambda]$ for $\boldsymbol{\alpha}$ and a flat prior for $\boldsymbol{\beta}$. The parameter λ, restricting α_i to a reasonably large interval, helps the algorithm converge faster but has little influence on the inference results, as shown in our simulation study. The tuning parameter σ_0 reflects the estimation of the relative dependence of rank information on covariates. The smaller σ_0 is, the more we rely on covariates in rank aggregation.

4.3 Probability interval for aggregated ranking list

In classical rank aggregation methods, people seek one aggregated rank only, but ignore the uncertainty of the aggregated result. When we notice $i \succ j$ in a single ranking list ρ, we cannot tell whether i is much better than j or they are close. Bayesian inference naturally provides us with an uncertainty measure of a ranking result. Suppose we have MCMC samples $\boldsymbol{\mu}^{(k)}, k = 1 \ldots, K$, from the distribution $p(\boldsymbol{\mu} \mid \tau_1^P, \cdots, \tau_m^P)$, where $\boldsymbol{\mu}^{(k)} = (\mu_1^{(k)}, \cdots, \mu_n^{(k)})'$. For each sample $\boldsymbol{\mu}^{(k)}$, we calculate a ranking list

$$\rho^{(k)} = \text{sort}(i \in U \text{ by } \mu_i^{(k)} \downarrow).$$

We denote $\tau^{(k)}(i)$ as the position of entity i in ranking list $\rho^{(k)}$. We define the $(1-\alpha)$ probability interval of entity i's rank as

$$\left(\tau^{LB}(i), \tau^{UB}(i)\right) = \left(\tau_{\lfloor(\frac{\alpha}{2})K\rceil}(i), \tau_{\lfloor(1-\frac{\alpha}{2})K\rceil}(i)\right),$$

where $\tau_{\lfloor(\frac{\alpha}{2})K\rceil}(i)$ and $\tau_{\lfloor(1-\frac{\alpha}{2})K\rceil}(i)$ are the $\lfloor(\frac{\alpha}{2})K\rceil$th and $\lfloor(1-\frac{\alpha}{2})K\rceil$th order statistics of $\{\tau^{(k)}(i)\}_{k=1,\ldots,K}$.

4.4 Extension to weighted rank aggregation

Considering that the rankers have different relative importance or reliability, people use weighted rank aggregation, in which larger weight w_j means that ranker j should play a more important role in the aggregated rank. However, it is difficult to assign a proper weight or design a proper weighting scheme in practice, especially when little or no prior knowledge of rankers is available. Under such circumstances, we handle the quality difference by incorporating weights as variance parameters in our model, and estimating them through Bayesian computation.

More precisely, we model the rankers' quality variations by allowing for different error variances, which extends model (6) to the weighted version as follows:

$$Z_{ij} = \alpha_i + x_i^T \boldsymbol{\beta} + \epsilon_{ij},\ \epsilon_{ij} \sim N(0, w_j^{-1}) \tag{7}$$

and

$$\tau_j^* = \tau(\boldsymbol{Z}_j),$$

where $\tau(z)$ is the full ranking list of z. Since the term $\alpha_i + x_i^T\boldsymbol{\beta}$ is regarded as the true value for entity i, the high quality rankers' evaluation should be closer to the truth, while the spam rankers' input lists should have large errors. We assume that the variance of error ϵ_{ij} depends on ranker j's quality, but does not depend on entity i. It is proportional to the inverse of the ranker's reliability, quantified as w_j.

We restrict w_j to three different levels: reliable ranker, mediocre ranker and spam ranker. The corresponding weights for these rankers are 2, 1 and 0.5, with equal probability as their prior. Explicitly, $p(\boldsymbol{w})$ is

$$P(w_j = 0.5) = P(w_j = 1) = P(w_j = 2) = \frac{1}{3}, \text{ for } \ j = 1, 2, \ldots, m; \tag{8}$$

and the w_js are mutually independent. Given this prior, the posterior distribution of $\boldsymbol{w}$ is

$$p(\boldsymbol{w} \mid \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m, \boldsymbol{\alpha}, \boldsymbol{\beta}) \propto p(\boldsymbol{w}) p(\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}),$$

which is a discrete distribution on $\{0.5, 1, 2\}$ with updated probability. We regard this weighted rank aggregation method as BARCW.

Besides the discrete candidate weights we just described, continuous weighting bounded away from 0 and ∞, which can be endowed with a chi-square prior, is another possible choice. Our numerical study shows that the more restricted candidate weights generate less autocorrelated MCMC samples in computation. Thus, discrete weights are more sensible to use in practice.

5 Bayesian computation with parameter expansion

We use Gibbs sampling with parameter expansion (Liu and Wu, 1999) in our Bayesian computation for the latent variable model with covariates. All models in Section 3 can be written in the following general form:

$$Z_{ij} \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w} \sim N(\alpha_i + x_i^T\boldsymbol{\beta}, w_j^{-1}) \text{ and } \tau_j^* = \tau(\boldsymbol{Z}_j).$$

Given the prior, we first consider the conditional posterior distribution of α_i:

$$\begin{aligned} & p(\alpha_i \mid \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m, \boldsymbol{\alpha}_{[-i]}, \boldsymbol{\beta}, \boldsymbol{w}, \tau_1^P, \cdots, \tau_m^P) \\ \propto \ & p(Z_{i,1}, \ldots, Z_{i,m} \mid \alpha_i, \boldsymbol{\beta}, \boldsymbol{w}) p(\alpha_i) \propto \exp\left(-\frac{\sum_{j=1}^m w_j (Z_{ij} - x_i^T\boldsymbol{\beta} - \alpha_i)^2}{2} - \frac{\alpha_i^2}{2\sigma_0^2} \right) 1_{[-\lambda, \lambda]}. \end{aligned}$$

The standard Gibbs sampler draws $[\alpha_i \mid \boldsymbol{\alpha}_{[-i]}, \boldsymbol{w}, \boldsymbol{\beta}, \boldsymbol{Z}]$ from the normal distribution $N(\hat{\mu}, v)$ truncated on $[-\lambda, \lambda]$, where $v = (\sum_j w_j + \sigma_0^{-2})^{-1}$ and $\hat{\mu} = \sum_j w_j(Z_{ij} - x_i^T\boldsymbol{\beta})/(\sum_j w_j + \sigma_0^{-2})$.

Second, we draw w_j, which is discrete, from the conditional posterior probability mass function:

$$\begin{aligned} & p(w_j \mid \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m, \boldsymbol{w}_{[-j]}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \tau_1^P, \cdots, \tau_m^P) \\ \propto \;& p(w_j)p(\boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m \mid \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{w}) \propto w_j^{\frac{n}{2}} \exp\left(-\frac{w_j \sum_{i=1}^n (Z_{ij} - x_i^T\boldsymbol{\beta} - \alpha_i)^2}{2}\right) \end{aligned}$$

Finally, we consider sampling $\boldsymbol{Z}$ and $\boldsymbol{\beta}$. The most computationally expensive part in our model is to sample all Z_{ij}s from the truncated normal. The relative scale of $\boldsymbol{\beta}$ can easily cause high correlation among $\boldsymbol{Z}$ and $\boldsymbol{\beta}$, similar to the classic Bayesian data augmentation method for probit regression (Albert and Chib, 1993). To speed up the algorithm, we conduct a group scale transformation on $\boldsymbol{Z}$. Conditioning on $\boldsymbol{\alpha}$ and $\boldsymbol{w}$, the complete data model can be expanded by introducing an expansion parameter θ for the residual's standard deviation, which is originally fixed at 1:

$$W_{ij} \mid \boldsymbol{\alpha}, \boldsymbol{\beta} \sim N(\alpha_i\theta + x_i^T\boldsymbol{\beta}\theta, \theta^2/w_j)$$

and $\tau_j^* = \tau(\boldsymbol{W}_j)$. This model can be derived from the scale transformation $\boldsymbol{Z} = t_\theta(\boldsymbol{W}) \equiv \boldsymbol{W}/\theta$. This parameter expansion is proper if we use the prior choice guided by the Haar measure (Liu and Wu, 1999). By assuming the Haar measure as the prior distribution for θ, $H(d\theta) = \theta^{-1}d\theta$, the Gibbs steps for $\boldsymbol{Z}$, θ and $\boldsymbol{\beta}$ are:

1. Draw $[W_{ij} \mid W_{[-i],j}, \boldsymbol{W}_{[-j]}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \theta]$ from $\mathcal{N}(\alpha_i\theta + x_i^T\boldsymbol{\beta}\theta, \theta^2/w_j)$ with truncation points determined by $W_{[-i],j}$, such that W_{ij} falls in the correct position according to τ_j^P. Then let $Z_{ij} = t_\theta(W_{ij}) = W_{ij}/\theta$.

2. Draw $\theta \sim p(\theta \mid \boldsymbol{Z}_1, \ldots, \boldsymbol{Z}_m, \boldsymbol{\alpha}, \tau_1^P, \cdots, \tau_m^P) \propto p\left(t_\theta(\boldsymbol{Z}) \mid \boldsymbol{\alpha}\right)|J_\theta(\boldsymbol{Z})\,|H(d\theta)$. Here, $|J_\theta(\boldsymbol{Z})| = \theta^{-nm}$ is the Jacobian of scale transformation, and

$$p(t_\theta(\boldsymbol{Z}) \mid \boldsymbol{\alpha}) \propto \int p(t_\theta(\boldsymbol{Z}) \mid \boldsymbol{\beta}, \boldsymbol{\alpha})p(\boldsymbol{\beta})d\boldsymbol{\beta} \propto \exp\left\{-\frac{S}{2\theta^2}\right\},$$

where

$$S = \sum_{j=1}^m w_j(\boldsymbol{Z}_j - \boldsymbol{\alpha})'(\boldsymbol{Z}_j - \boldsymbol{\alpha}) - \frac{1}{\sum_{j=1}^m w_j}\sum_{j,k}^m w_j w_k(\boldsymbol{Z}_j - \boldsymbol{\alpha})'X(X'X)^{-1}X'(\boldsymbol{Z}_k - \boldsymbol{\alpha}).$$

Here $X = (x_1, x_2, \ldots, x_n)$ is the matrix whose columns are the covariates vectors for each entity. Thus we have $\theta^2 \sim S/\chi^2_{nm}$.

3. Draw $\boldsymbol{\beta} \sim p(\boldsymbol{\beta} \mid \boldsymbol{\alpha}, t_\theta(\boldsymbol{Z})) \equiv N\left(\hat{\boldsymbol{\beta}}/\theta, V\right)$, where

$$\hat{\boldsymbol{\beta}} = \frac{1}{\sum_{j=1}^m w_j}(X'X)^{-1}X'\left[\sum_{j=1}^m w_j(\boldsymbol{Z}_j - \boldsymbol{\alpha})\right]$$

and $V = \frac{1}{\sum_{j=1}^m w_j}(X'X)^{-1}$.

The above sampling scheme is the implementation of Parameter Expansion for Data Augmentation (PX-DA) (Liu and Wu, 1999) for our model. After we obtain Monte Carlo samples from the posterior distribution $p(\alpha_i + x_i^T\boldsymbol{\beta} \mid \tau_1^P, \cdots, \tau_m^P)$, we use the posterior mean of $\alpha_i + x_i^T\boldsymbol{\beta}$'s to generate the aggregated ranking list. The samples of $p(w_i \mid \tau_1^P, \cdots, \tau_m^P)$ measures the reliability of the input rankers. If we wish to assign equal weight or predetermined unequal weight to the rankers, we can fix $\boldsymbol{w}$ in our MCMC computation.

6 Simulation Studies

6.1 Simulation under the BARC model

Let U be the set $\{1, 2, \ldots, n\}$, where entity i has true value μ_i. We generated m full ranking lists $\{\tau_j\}_{1 \le j \le m}$ via the following scheme:

$$\tau_j = \text{sort}(i \in U \text{ by } Z_{ij} \downarrow), \text{ where } Z_{ij} \sim \mathcal{N}(\mu_i, \sigma^2).$$

We generated $p \times 1$ covariates vector x_i from an i.i.d. multivariate normal distribution with mean 0 and covariance $Cov(x_{is}, x_{it}) = \rho^{|s-t|}, 1 \le t \le p$. We examined three scenarios. In Scenario 1, the true difference between entities can be linearly explained by covariates. In Scenario 2, a linear combination of covariates can partially explain the ranking. In Scenario 3, the ranking mechanism is barely correlated with covariates.

1. $\mu_i = x_i^T\boldsymbol{\beta}$, with coefficient $\boldsymbol{\beta} = (3, 2, 1, 0.5)'$, $\rho = 0.2$.

2. $\mu_i = x_i^T\boldsymbol{\beta} + \|x_i\|^2$, with coefficient $\boldsymbol{\beta} = (3, 2, 1)'$, $\rho = 0.5$.

3. $\mu_i = \|x_i\|^2$, $\rho = 0.5$.

We first examined the impact of noise level σ on the performance of the BARC algorithm and other rank aggregation methods. Fixing $n = 50$ and $m = 10$, we tried five different values of σ ($\sigma = 1, 2, 5, 10, 20$). For each configuration, we generated 500 simulated datasets. We applied three naive methods (AriM, GeoM, MedR), three Markov-Chain based methods (MC1, MC2, MC3) and our BARC methods with four different pairs of hyper-parameters (λ, σ_0^2). The distance between true rank and the aggregated rank produced by these methods in each simulated data were used to evaluate the performances of different methods. Figure 1 plots the average Kendall's tau distance against noise level over 500 simulations for the methods mentioned above. Spearman's footrule distance gives a very similar result and is hence omitted here.

From Figure 1, we can see that: (1) when noise level is very small, all methods performed similarly well in recovering the true rank; (2) BARC uniformly outperformed the competing methods in all three scenarios; (3) the difference between BARC and the other methods is larger in Scenarios 1 and 2, where the linear combination of covariates is more useful; (4) BARC is robust with all four hyper-parameter choices—the BARC lines with different hyper-parameters λ (=5,10) overlapped, while the BARC lines with different hyper-parameters σ_0^2 (=0.5,1) were very close to each other.

Next, we evaluate the influence of sample sizes (i.e. the number of entities n and the number of ranking lists m) on the performances of the rank aggregation methods. To figure out the impact of the number of rankers m, we fixed the noise level at $\sigma = 5$, the number of entities at $n = 50$, and let m take values 5, 10, 20 and 40 respectively. Likewise, we also explored the impact of n by fixing $m = 10$ and $\sigma = 5$. The results are summarized in Figure 2, from which we can see that BARC is quite robust to different data sizes.

6.2 Partial ranking lists

We further explore how BARC works when aggregating partial ranking lists, where subgroups have no overlap with each other. We simulated data from Scenario 2 with $n = 80$, $m = 10$. We randomly divide these 80 entities into $k\,(= 1, 2, 4, 8, 10, 16)$ subgroups, each with size

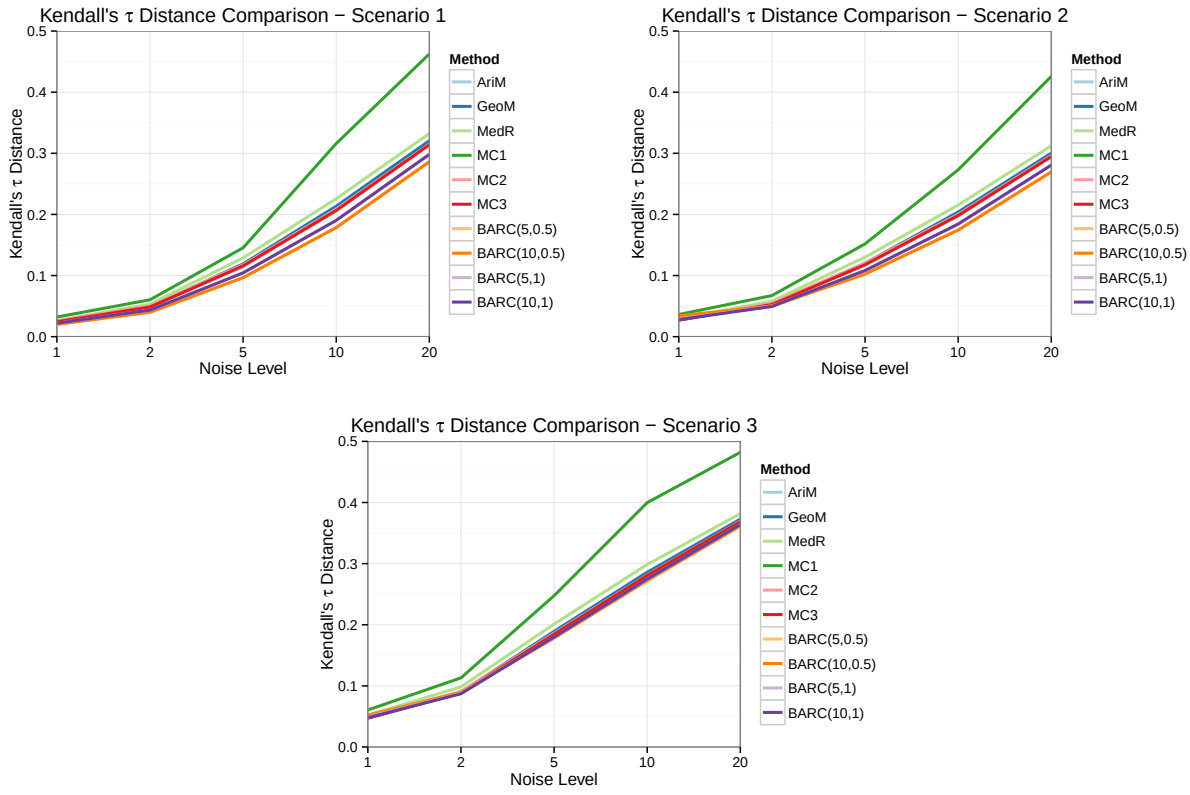

Figure 1: Average distance between true rank and aggregated rank of different methods of rank aggregation. As the ranking mechanism is decreasingly correlated with covariates from Scenario 1 to Scenario 3, the advantage of BARC over existing methods shrinks.

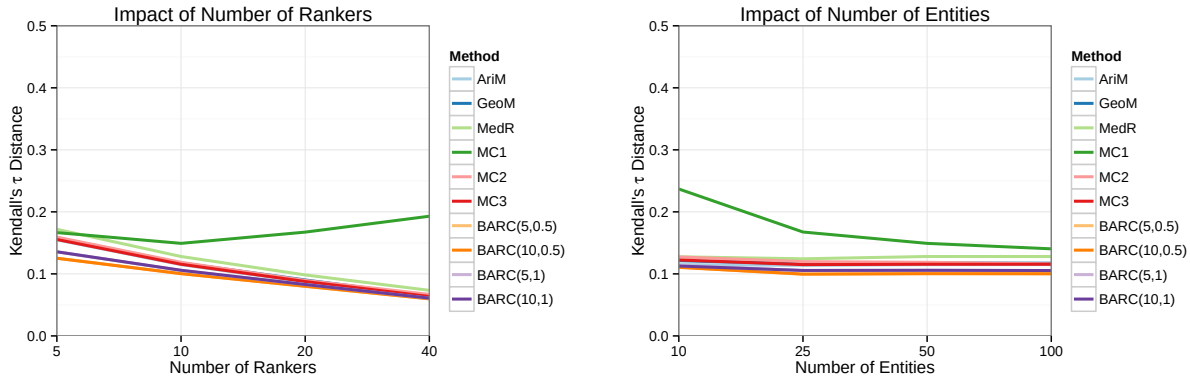

Figure 2: Impact of data size on the performance of different methods.

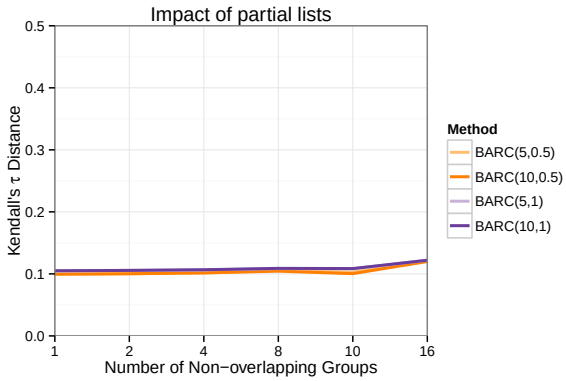

Figure 3: Robustness of BARC with respect to partial ranking lists. With the help of covariates, BARC achieve nearly the same performance level as the incompleteness of the ranking lists increase.

m/k. As k increases, the pairwise comparison information decreases. For example, when $k = 16$, we have only $16 \times \binom{5}{2} = 160$ pairwise comparisons in a partial ranking list, while a full ranking list has $\binom{80}{2}$. Figure 3 displays the Kendall tau distance from the aggregated ranking list calculated by BARC to the true rank in different cases. The BARC is quite robust with respect to partial ranking lists when unobserved pairwise comparisons are missing completely at random and the input ranking lists depend on covariates.

6.3 Robustness to spam rankers

This subsection evaluates the performance of BARC and BARCW in the context of a difference in quality among the input lists. In particular, we consider a top-d list rank aggregation similar to the example of five microarray studies of prostate cancer in DeConde et al. (2006). We generated $m = 5$ input lists $\{\tau_j\}_{1 \le j \le m}$, each with $n = 200$ items, from the following model:

$$\tau_j = \text{sort}(i \in U \text{ by } Z_{ij} \downarrow), \text{ where } Z_{ij} \sim \mathcal{N}(\mu_i, \sigma_j^2),$$

$$\mu_i = x_i^T \boldsymbol{\beta}, \text{ with coefficient } \boldsymbol{\beta} = (3, 2, 1)'.$$

Among the 5 input lists, we assigned them different levels of quality. If τ_j is from a reliable ranker, then $\sigma_j = 5$; if τ_j is from a spam ranker, then $\sigma_j = 200$. Spam rankers provide highly random ranking lists. We tried 5 different numbers of reliable ranking lists. Using 200

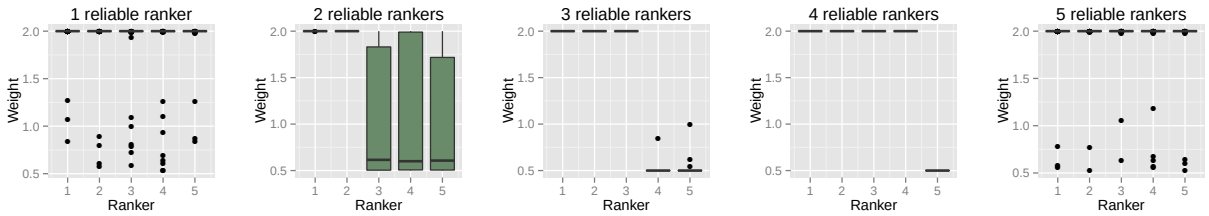

Figure 4: Boxplots of posterior means of weights given by BARCW in 200 simulations. The dots are considered outliers outside of the range between first and third quartile of the 200 posterior means of weights. A black line, which is indeed a collapsed box, at level 2.0 indicates that the minimum of the posterior mean is almost 2, so the corresponding ranker is identified by BARCW as a reliable ranker. From left to right, the number of reliable rankers increases.

simulations each, we examined the performance of different methods when different numbers of reliable input lists exist. Posterior means of the weights by the BARCW method are shown as boxplots in Figure 4. In addition to the comparison of distance measure, we also compare the true positive rate of the methods to top-d list aggregation ($d = 10$).

As shown in Figure 4, BARCW almost assigned equal weights for all 5 rankers in the first and last scenarios. When there was only one reliable ranker, the reliable ranker was not identifiable. In the middle three scenarios, BARCW successfully identified the reliable rankers, with only a few spam rankers misspecified in the situation where two reliable rankers exist.

BARCW gave the best result in the middle three situations, but failed to beat BARC when the reliable list is non-identifiable or all input lists are equally reliable (see Figure 5). BARC uniformly outperformed the other methods in this simulation, where the ranking scheme depends on covariates, even though the input lists have different reliabilities.

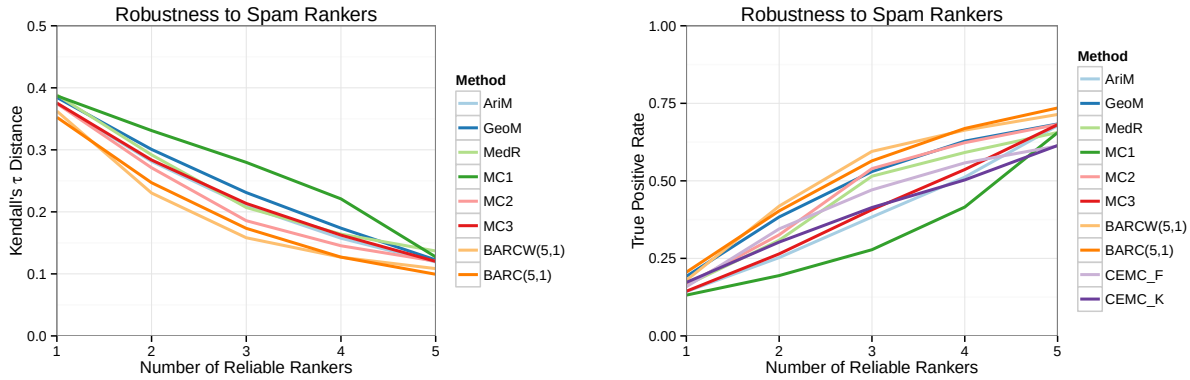

Figure 5: Kendall's τ distance and true positive rate of different methods. CEMC_F and CEMC_K stand for cross entropy Monte Carlo approaches. They are only included in the top-d list comparison because they are computationally too expensive to aggregate a long full list.

7 Real Data Applications

7.1 Aggregating NFL Quarterback

The NFL quarterback ranking is a classic case where experts' ranking scheme is clearly related to players' statistics in their games. Using Table 1 and Table 2, we generated aggregated lists using both rank data and covariates information. Table 4 shows the results of major rank aggregation methods. For quarterbacks at the top and bottom of the list, these methods mostly agree with each other. Figure 6 shows the 95% probability interval for each quarterback's rank. We can see that the interval width is large for mediocre quarterbacks, and that's exactly where there is the most discrepancy among different rankers and different rank aggregation methods.

All methods except BARCW assume equal reliability for all input lists. Figure 7 shows the boxplots and means of posterior samples of weights. Out of 13 rankers, seven are learned to have significantly higher quality than the other six rankers. In this application, BARCW tends to reinforce the majority's opinion, which provided similar rankings. It can sometimes be harmful—for example, when a few reliable rankers have very different opinions than the other high quality rankers.

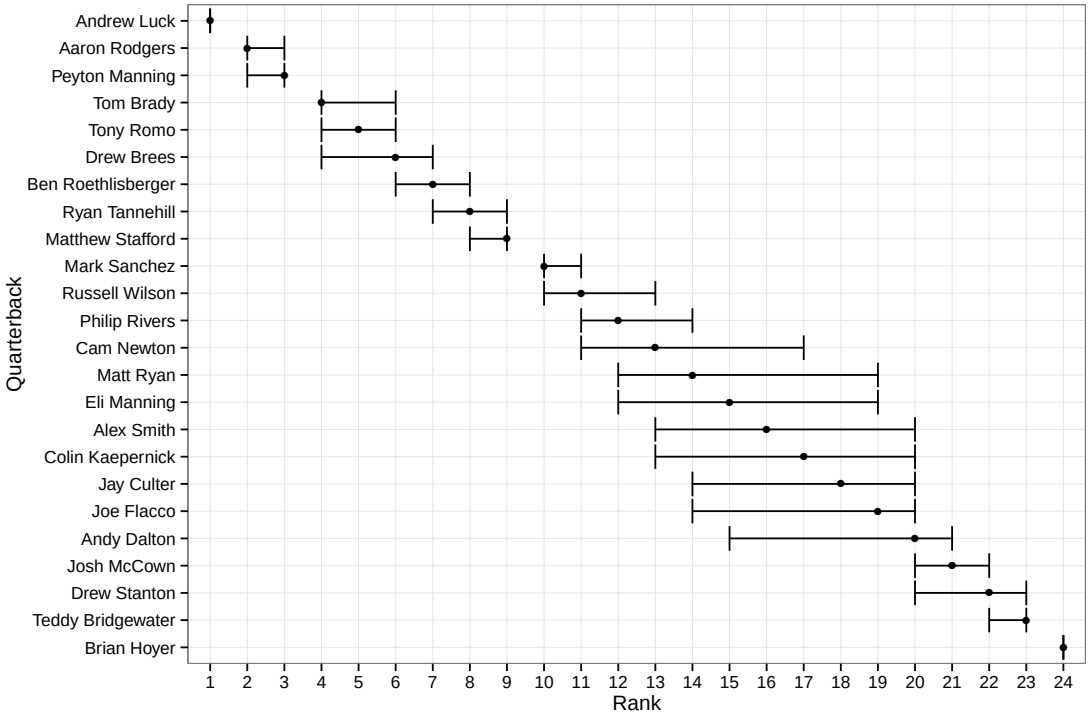

Figure 6: Interval estimates of aggregated ranks of NFL quarterbacks, as of week 12 in the 2014 NFL season.

Table 4: Rank aggregation results of NFL quarterbacks listed in the order of BARCW posterior means. The rankings are listed to the right of the underlying values given by rank aggregation methods.

Player	BARCW μ	Rank	BARC μ	Rank	AriM	Rank	MC1.π	Rank	MC2.π	Rank	MC3.π	Rank
Andrew Luck	6.518	1	6.069	1	1.286	1	0.143	1	0.225	1	0.207	1
Aaron Rodgers	4.760	2	4.635	2	2.571	2	0.115	2	0.153	2	0.137	2
Peyton Manning	4.466	3	4.39	3	3	3	0.105	3	0.11	3	0.12	3
Tom Brady	2.937	4	2.942	4	5.071	4	0.078	5	0.084	4	0.071	4
Tony Romo	2.805	5	2.744	5	5.214	5	0.083	4	0.066	5	0.07	5
Drew Brees	2.469	6	2.448	6	5.857	6	0.073	6	0.048	7	0.063	6
Ben Roethlisberger	2.149	7	2.094	7	6.571	7	0.06	7	0.048	6	0.052	7
Ryan Tannehill	1.435	8	1.342	8	8	8	0.042	9	0.036	8	0.04	8
Matthew Stafford	0.965	9	0.72	9	8.857	9	0.042	8	0.031	9	0.034	9
Mark Sanchez	-0.005	10	-0.098	10	11.5	10	0.022	11	0.026	10	0.023	10
Russell Wilson	-0.716	11	-0.496	11	12.214	11	0.021	12	0.023	11	0.021	11
Philip Rivers	-0.930	12	-0.602	12	13.214	12	0.028	10	0.019	12	0.02	12
Cam Newton	-1.197	13	-1.136	13	14.5	13	0.019	13	0.017	13	0.017	13
Matt Ryan	-1.413	14	-1.474	14	16.357	15	0.015	20	0.014	15	0.014	15
Eli Manning	-1.474	15	-1.497	15	16.071	14	0.016	17	0.017	14	0.014	14
Alex Smith	-1.793	16	-1.717	17	16.714	16	0.016	18	0.011	19	0.013	16
Colin Kaepernick	-1.813	17	-1.601	16	16.786	17	0.018	15	0.011	18	0.013	17
Joe Flacco	-1.815	18	-1.778	19	16.929	18	0.016	16	0.012	16	0.013	18
Jay Culter	-1.884	19	-1.726	18	17.143	19	0.019	14	0.01	20	0.013	19
Andy Dalton	-1.987	20	-2.052	20	17.357	20	0.016	19	0.011	17	0.012	20
Josh McCown	-2.733	21	-2.645	21	19.5	21	0.015	22	0.008	21	0.01	21
Drew Stanton	-2.812	22	-2.823	22	20.286	22	0.015	21	0.008	22	0.009	22
Teddy Bridgewater	-3.476	23	-3.378	23	21.714	23	0.013	23	0.007	23	0.008	23
Brian Hoyer	-4.462	24	-4.361	24	23.286	24	0.009	24	0.006	24	0.007	24

Fortunately, we can validate our weights estimation in NFL quarterback ranking data using the prediction accuracy of each expert at the end of the season. *FantasyPros.com* rates the accuracy of each expert's ranking by comparing their player predictions to the actual outcomes. Experts are awarded for correct predictions based on the incremental fantasy points achieved by making correct lineup decisions. Figure 8 plots the posterior means of weights w versus each expert's accuracy. Rankers with higher weights mostly outperform those with lower weights.

In BARC and BARCW, models are over-parameterized on β and α. Although their values are not very meaningful, Figure 9 can still give us some intuition about the role of covariates in our rank aggregation. TD% and Int%, which stand for percentage of touchdowns and interceptions thrown when attempting to pass, are the most significant covariates; touchdowns have a positive effect, while interceptions is have a negative one. Based on our football common sense, touchdowns and interceptions can directly impact the result of a game.

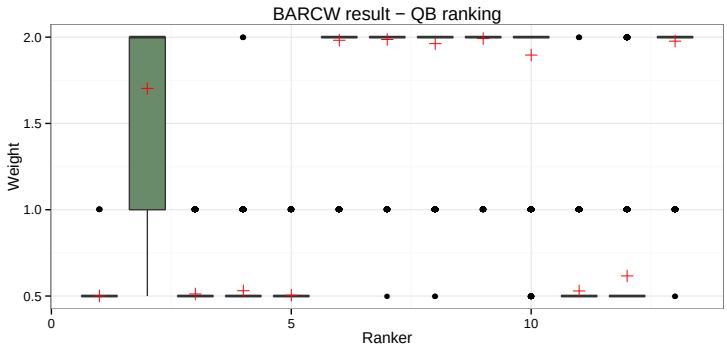

Figure 7: Boxplots of posterior samples of weights given by BARCW. Red cross marks the posterior means of weights. Black points are samples outside of the range between first and third quartile of the posterior samples, and black lines are collapsed boxes when interquartile range is 0. Seven experts are learned to be reliable rankers.

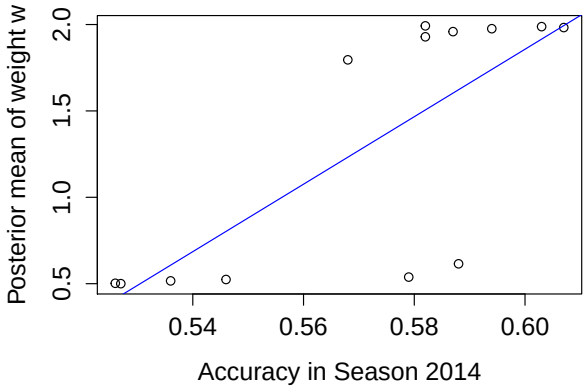

Figure 8: Scatter plot of the posterior means of weights versus experts' accuracy evaluated by *FantasyPros.com* after the 2014 NFL season. The blue line is a fitted least-square line. The positive correlation between weights and accuracy is significant with a p-value of 0.003.

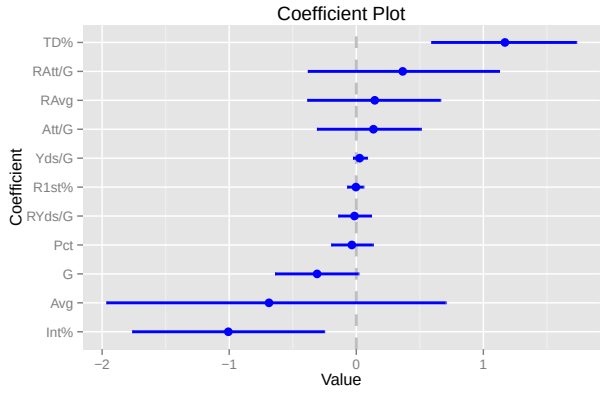

Figure 9: Posterior mean and 95% probability interval of β given by BARCW in aggregating quarterback rankings.

Covariates play a more important role in cases where input lists are incomplete. We split our entities into two groups in three different ways:

Case 1: We randomly divide these quarterbacks into two groups.

Case 2: Split quarterbacks into two groups based on Borda Count.

Case 3: Split quarterbacks into two groups based on American Football Conference (AFC) and National Football Conference (NFC).

Table 5 shows the BARC aggregation results in the above cases. In Case 1 and Case 3, the two groups of quarterbacks were on generally the same level, so BARC generated similar ranking lists to the one from full list aggregation. In Case 2, one group of quarterbacks are better than the others. Benefitting from the usage of covariates, BARC figured out that one group is better than the other and provided a reasonable list, whose top eight are the same as the full list aggregation.

7.2 Aggregating orthodontics data

As mentioned in Section 1, the orthodontics data set contains 69 lists for each of the 9 groups of patients. Table 3 shows that the rankers did not have very similar opinions. The covariates for these patients are objective measurements about their teeth. We applied both BARC

Table 5: Rank aggregation results of combining partial lists created by splitting quarterback rankings in three different ways.

Player	BARC.full.list.	Rank	BARC.Case.1	Rank.1	BARC.Case.2	Rank.2	BARC.Case.3	Rank.3
Andrew Luck	6.043	1	5.803	1	4.823	1	5.870	1
Aaron Rodgers	4.612	2	4.637	2	3.407	2	4.580	2
Peyton Manning	4.360	3	3.735	3	3.026	3	4.051	3
Tom Brady	2.930	4	2.713	4	1.752	4	2.638	6
Tony Romo	2.726	5	2.501	5	1.567	5	3.001	4
Drew Brees	2.434	6	2.083	6	1.279	6	2.756	5
Ben Roethlisberger	2.086	7	1.661	7	0.895	7	1.613	7
Ryan Tannehill	1.347	8	1.000	8	0.286	8	0.955	9
Matthew Stafford	0.727	9	0.189	9	-0.257	11	1.174	8
Mark Sanchez	-0.079	10	-0.635	12	-1.178	18	0.224	10
Russell Wilson	-0.477	11	0.075	10	-1.885	22	-0.258	11
Philip Rivers	-0.579	12	-0.221	11	-1.740	21	-1.049	12
Cam Newton	-1.111	13	-0.743	13	0.068	9	-1.375	16
Matt Ryan	-1.467	14	-1.055	15	-0.282	12	-1.189	13
Eli Manning	-1.493	15	-1.029	14	-0.191	10	-1.300	15
Colin Kaepernick	-1.599	16	-1.823	18	-0.442	13	-1.223	14
Alex Smith	-1.704	17	-1.245	16	-0.529	14	-1.969	18
Jay Culter	-1.715	18	-1.896	19	-0.583	16	-1.446	17
Joe Flacco	-1.770	19	-1.332	17	-0.580	15	-1.992	19
Andy Dalton	-2.040	20	-2.322	21	-0.824	17	-2.150	20
Josh McCown	-2.640	21	-2.140	20	-1.463	19	-2.768	22
Drew Stanton	-2.845	22	-2.532	22	-1.723	20	-2.453	21
Teddy Bridgewater	-3.377	23	-3.451	23	-2.210	23	-3.173	23
Brian Hoyer	-4.361	24	-3.984	24	-3.226	24	-4.509	24

and BARCW to this data set. Table 6 shows the top 20 of 108 patients in our aggregations.

Figure 10 shows the posterior weights of rankers. Roughly half of the experts are assigned weight 0.5. Due to the difference in weights, the results of BARC and BARCW are quite different in this application. The rank aggregation of subgroup ranking lists is hard, and no existing methods can combine non-overlapping ranking lists.

8 Discussion

In this article, we propose two Bayesian rank aggregation methods (BARC and BARCW) for the rank aggregation problem when covariates of entities are also observed. BARC and BARCW use similar latent variable models to incorporate covariates information. With the help of covariates, BARC and BARCW can accommodate various types of input ranking lists, including highly incomplete ones. Moreover, BARCW learns the qualities of rankers from data, and naturally over-weighs high quality ones in rank aggregation. Furthermore,

Table 6: Top 20 of 108 entities given by BARC and BARCW in aggregating Orthodontics data.

	Entity	BARC.μ	Entity	BARCW.μ
1	E2	6.559	G7	2.638
2	G7	6.531	E2	2.565
3	C7	6.042	H2	1.79
4	C6	5.95	F8	1.608
5	C1	5.808	E1	1.573
6	C5	5.78	H3	1.513
7	A1	5.688	E10	1.495
8	E1	5.674	G1	1.477
9	E10	5.672	E4	1.444
10	G1	5.552	H4	1.393
11	D2	5.538	A1	1.386
12	H2	5.518	F2	1.289
13	F8	5.5	D2	1.223
14	C4	5.483	F1	1.22
15	E4	5.421	B3	1.203
16	A4	5.405	A4	1.176
17	I2	5.383	A8	0.999
18	B3	5.308	H6	0.987
19	A8	5.303	I2	0.955
20	C12	5.275	G2	0.941

our Bayesian inference generates uncertainty measures for the aggregated ranks, which is not provided by existing rank aggregation methods. In simulation studies, we showed that both BARC and BARCW outperformed the existing methods in different scenarios. Our real-data application validates our estimation of reliability parameters and effects of covariates.

In general, BARC and BARCW are robust to different choices of hyper-parameters, σ_0 and λ. In practice, we may need to try different hyper-parameter settings to ensure the robustness before a conclusion can be made. Due to the difference in the size of rank-data, a MCMC convergence diagnostic is always needed.

Although BARCW directly quantifies the difference in qualities of rankers, we have limited abilities to define qualities by the nature of our model setting. Rankers' reliability in real world can be highly complex. For example, there can be a few groups of rankers with different opinions, despite all being reliable rankers. In this case, it is desirable to learn the structure of rankers instead of simply classifying some of them as spam rankers.

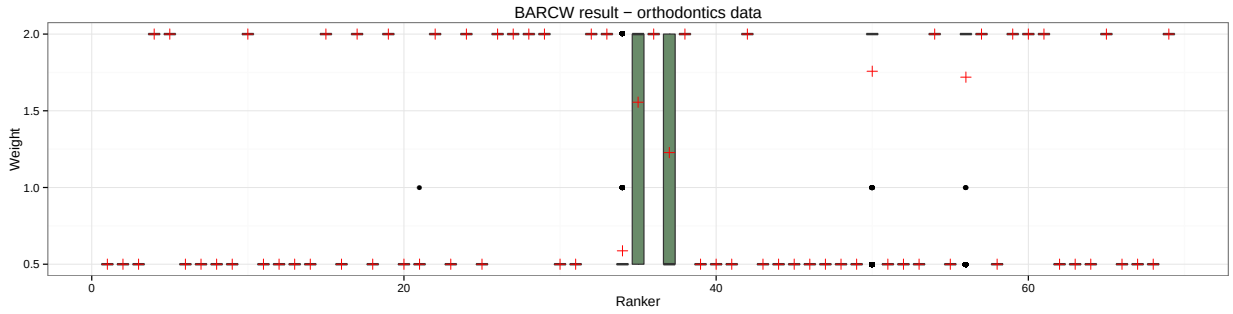

Figure 10: Boxplot of posterior samples of weights determined by BARCW in aggregating orthodontics data. Red cross marks the posterior means of weights. Black points are samples outside of the range between first and third quartile of the posterior samples, and black lines are collapsed boxes when interquartile range is 0.

References

James H. Albert and Siddhartha Chib. Bayesian analysis of binary and polychotomous response data. *Journal of the American statistical Association*, 88:669–679, 1993.

Jean C. de Borda. Mémoire sur les élections au scrutin. 1781.

Robert P. DeConde, Sarah Hawley, Seth Falcon, Nigel Clegg, Beatrice Knudsen, and Ruth Etzioni. Combining results of microarray experiments: a rank aggregation approach. *Statistical Applications in Genetics and Molecular Biology*, 5, 2006.

Ke Deng, Simeng Han, Kate J. Li, and Jun S. Liu. Bayesian aggregation of order-based rank data. *Journal of the American Statistical Association*, 109:1023–1039, 2014.

Persi Diaconis and Ronald L. Graham. Spearman's footrule as a measure of disarray. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 262–268, 1977.

Cynthia Dwork, Ravi Kumar, Moni Naor, and Dandapani Sivakumar. Rank aggregation methods for the web. In *Proceedings of the 10th international conference on World Wide Web*, pages 613–622. ACM, 2001.

Maurice G. Kendall. A new measure of rank correlation. *Biometrika*, pages 81–93, 1938.

Shili Lin. Rank aggregation methods. *Wiley Interdisciplinary Reviews: Computational Statistics*, 2:555–570, 2010.

Shili Lin and Jie Ding. Integration of ranked lists via cross entropy monte carlo with applications to mrna and microrna studies. *Biometrics*, 65:9–18, 2009.

Jun S. Liu and Ying Nian Wu. Parameter expansion for data augmentation. *Journal of the American Statistical Association*, 94:1264–1274, 1999.

Albert Maydeu-Olivares. Thurstonian modeling of ranking data via mean and covariance structure analysis. *Psychometrika*, 64:325–340, 1999.

Reuven Y. Rubinstein and Dirk P. Kroese. *The cross-entropy method: a unified approach to combinatorial optimization, Monte-Carlo simulation and machine learning*. Springer Science & Business Media, 2004.

Martin A. Tanner and Wing Hung Wong. The calculation of posterior distributions by data augmentation. *Journal of the American statistical Association*, 82:528–540, 1987.

Louis L. Thurstone. A law of comparative judgment. *Psychological review*, 34:273, 1927.

Aad W. Van der Vaart. *Asymptotic statistics*, volume 3. Cambridge university press, 2000.

Grace Yao and Ulf Böckenholt. Bayesian estimation of thurstonian ranking models based on the gibbs sampler. *British Journal of Mathematical and Statistical Psychology*, 52:79–92, 1999.

H. Peyton Young. Condorcet's theory of voting. *American Political Science Review*, 82: 1231–1244, 1988.

Appendix: Proof for the consistency of MLE

We need the following two lemmas to prove Theorem 1.

Lemma 1 *Given $\boldsymbol{\mu} \in \boldsymbol{\Theta}$, for every $\epsilon > 0$, there exists R such that for all $\boldsymbol{v} \in \boldsymbol{\Theta}$ with $\|\boldsymbol{\mu} - \boldsymbol{v}\| < 1$,*

$$\int_{\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} < \epsilon$$

where $\mathcal{B}(R) = \{\boldsymbol{Z} \in \mathbb{R}^n : \ \|\boldsymbol{Z}\| \leq R\}$.

Proof of lemma 1. According to the triangle inequality and the fact that $\|\boldsymbol{v} - \boldsymbol{\mu}\| < 1$,

$$\|\boldsymbol{Z} - \boldsymbol{v}\| \geq \|\boldsymbol{Z}\| - \|\boldsymbol{v} - \boldsymbol{u}\| - \|\boldsymbol{u}\| \geq \|\boldsymbol{Z}\| - 1 - \|\boldsymbol{u}\| = \|\boldsymbol{Z}\|/2 + (\|\boldsymbol{Z}\|/2 - 1 - \|\boldsymbol{u}\|)\,.$$

Therefore, when $R \geq 2(1 + \|\boldsymbol{\mu}\|)$, $\|\boldsymbol{Z} - \boldsymbol{v}\| \geq \|\boldsymbol{Z}\|/2$ for any $\boldsymbol{Z} \in \mathcal{B}(R)^c$, and thus

$$\int_{\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \leq \int_{\mathcal{B}(R)^c} e^{-\frac{1}{8}\|\boldsymbol{Z}\|^2} d\boldsymbol{Z}. \tag{9}$$

Note that the right hand side of inequality (9) converges to zero as R goes to infinity. For any $\epsilon > 0$, there exists R such that $R \geq 2(1 + \|\boldsymbol{\mu}\|)$ and $\int_{\mathcal{B}(R)^c} e^{-\frac{1}{8}\|\boldsymbol{Z}\|^2} d\boldsymbol{Z} < \epsilon$, which further implies $\int_{\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} < \epsilon$ for any $\|\boldsymbol{v} - \boldsymbol{\mu}\| < 1$. □

Lemma 2 *(Identifiability) The true distribution of τ is identifiable on parameter space $\boldsymbol{\Theta}$: $P_{\boldsymbol{\mu}} \neq P_{\boldsymbol{\mu}'}$ for every $\boldsymbol{\mu} \neq \boldsymbol{\mu}'$.*

Proof of lemma 2. Suppose $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \boldsymbol{\Theta}$ and $P_{\boldsymbol{\mu}}(\tau) = P_{\boldsymbol{\mu}'}(\tau)$ for all possible full ranking list τ on U. Since

$$P_{\boldsymbol{\mu}}(\tau(i_1) < \tau(i_2)) = \sum_{\tau(i_1)<\tau(i_2)} P_{\boldsymbol{\mu}}(\tau),$$

we have $P_{\boldsymbol{\mu}}(\tau(i_1) < \tau(i_2)) = P_{\boldsymbol{\mu}'}(\tau(i_1) < \tau(i_2))$ for any $i_1, i_2 \in U$. On the other hand, $P_{\boldsymbol{\mu}}(\tau(i_1) < \tau(i_2)) = \Phi(\frac{\mu_{i_1} - \mu_{i_2}}{\sqrt{2}})$, where $\Phi(\cdot)$ is the Normal CDF. Hence, $\mu_{i_1} - \mu_{i_2} = \mu'_{i_1} - \mu'_{i_2}$ for $1 \leq i_1 < i_2 \leq n$, and $\mathbf{1}'\boldsymbol{\mu} = \mathbf{1}'\boldsymbol{\mu}' = 0$. Hence, we have $\boldsymbol{\mu} = \boldsymbol{\mu}'$. □

Proof of Theorem 1. We apply Wald's consistency proof (Van der Vaart, 2000) by verifying the following three conditions.

First, we denote $m_{\boldsymbol{\mu}}(\tau) \stackrel{def}{=} \log P_{\boldsymbol{\mu}}(\tau)$ and show that the map $\boldsymbol{\mu} \mapsto m_{\boldsymbol{\mu}}(\tau)$ is continuous for all τ. According to Lemma 1, given $\boldsymbol{\mu} \in \boldsymbol{\Theta}$, for every $\epsilon > 0$, there exists R such that

for all $\boldsymbol{v} \in \boldsymbol{\Theta}$ with $\|\boldsymbol{\mu} - \boldsymbol{v}\| < 1$, $\int_{\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} < \epsilon/4$. Thus, for any $\mathcal{A} \subset \mathbb{R}^n$ and $\|\boldsymbol{v} - \boldsymbol{\mu}\| < 1$,

$$\int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} + \int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} < \epsilon/2.$$

Since $e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2}$ is bound by 1 and $\mathcal{A} \cap \mathcal{B}(R)$ is a bounded area, we can apply Bounded Convergence Theorem: for any sequence $\{\boldsymbol{\mu}_k\}_{k=1}^{\infty}$ converging to $\boldsymbol{\mu}$,

$$\lim_{k\to\infty} \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}_k\|^2} d\boldsymbol{Z} = \int_{\mathcal{A}\cap\mathcal{B}(R)} \lim_{k\to\infty} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}_k\|^2} d\boldsymbol{Z} = \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z}.$$

Hence, $\int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z}$ is a continuous function of $\boldsymbol{\mu}$, and thus for every $\epsilon > 0$, there exist δ such that for all $\boldsymbol{v} \in \boldsymbol{\Theta}$ with $\|\boldsymbol{\mu} - \boldsymbol{v}\| < \delta$,

$$\left| \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} - \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \right| \leq \epsilon/2$$

Therefore, given $\boldsymbol{\mu} \in \boldsymbol{\Theta}$ and $\mathcal{A} \subset \mathbb{R}^n$, for every $\epsilon > 0$, there exists R such that for all $\boldsymbol{v} \in \boldsymbol{\Theta}$ with $\|\boldsymbol{\mu} - \boldsymbol{v}\| < \min\{1, \delta\}$,

$$\begin{aligned}
& \left| \int_{\mathcal{A}} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} - \int_{\mathcal{A}} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \right| \\
\leq & \left| \int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} - \int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \right| + \left| \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} - \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \right| \\
\leq & \int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} + \int_{\mathcal{A}\cap\mathcal{B}(R)^c} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} + \left| \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z} - \int_{\mathcal{A}\cap\mathcal{B}(R)} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{v}\|^2} d\boldsymbol{Z} \right| \\
< & \epsilon.
\end{aligned}$$

We conclude that $\int_{\tau(\boldsymbol{Z})=\tau'} e^{-\frac{1}{2}\|\boldsymbol{Z}-\boldsymbol{\mu}\|^2} d\boldsymbol{Z}$ is a continuous function with respect to $\boldsymbol{\mu}$.

Second, for every sufficiently small $V \subset \boldsymbol{\Theta}$ the function $\tau \mapsto \sup_{\boldsymbol{\mu}\in U} m_{\boldsymbol{\mu}}(\tau)$ is measurable and satisfies

$$E_{\boldsymbol{\mu}_0} \sup_{\boldsymbol{\mu}\in V} m_{\boldsymbol{\mu}}(\tau) < \infty.$$

The domain of τ a finite set containing all possible ranking lists of n entities. Thus, for any $t \in R$, the preimage of (t, ∞) under $\tau \mapsto \sup_{\boldsymbol{\mu}\in U} m_{\boldsymbol{\mu}}(\tau)$ is a finite set. Thus, $\tau \mapsto \sup_{\boldsymbol{\mu}\in U} m_{\boldsymbol{\mu}}(\tau)$ is a measureable function.

Third, because $P_{\boldsymbol{\mu}}(\tau) \leq 1$, we have $\sup_{\boldsymbol{\mu}\in V} m_{\boldsymbol{\mu}}(\tau) \leq 0$. Since the domain of τ is a finite set, there exists a lower bound c such that $\sup_{\boldsymbol{\mu}\in V} m_{\boldsymbol{\mu}}(\tau) \geq c$ for every τ. Thus, $E_{\boldsymbol{\mu}_0} \sup_{\boldsymbol{\mu}\in V} m_{\boldsymbol{\mu}}(\tau)$ exists and $E_{\boldsymbol{\mu}_0} \sup_{\boldsymbol{\mu}\in V} m_{\boldsymbol{\mu}}(\tau) < \infty$.

Due to the identifiability of $P_{\boldsymbol{\mu}}$ in Lemma 2, $E_{\boldsymbol{\mu}_0} m_{\boldsymbol{\mu}}(\tau)$ attains its maximum uniquely at $\boldsymbol{\mu}_0$. Then according to Wald's consistency proof, for every $\epsilon > 0$ and every compact set $K \subset \boldsymbol{\Theta}$,

$$P\left(\{\|\hat{\boldsymbol{\mu}}_m - \boldsymbol{\mu}_0\| \geq \epsilon\} \cap \{\hat{\boldsymbol{\mu}}_m \in K\}\right) \to 0, \text{ as } n \text{ is fixed}, m \to \infty.$$